

DROPPED DATABASE TABLE RECOVERY

RELATED APPLICATIONS

This patent application claims priority from the commonly assigned Canadian
5 Patent Application entitled "Dropped Database Table Recovery", having Canadian Patent
Application Serial No. 2,279,028, filed on July 29, 1999 by Effi Offer, Roger Luo Q.
Zheng, Matthew A. Huras, Michael J. Winer, and Dale M. McInnis and is a continuation
of and claims the benefit of a U.S. Patent Application entitled "Dropped Database Table
Recovery," having U.S. Application No. 09/626,673, filed on July 27, 2000 by Effi Offer,
10 Roger Luo Q. Zheng, Matthew A. Huras, Michael J. Winer, and Dale M. McInnis, each
of which applications are incorporated herein by reference in their entirety.

1. Field of the Invention

The present invention is directed to an improvement in database systems and in
15 particular to the recovery of tables dropped from databases.

2. Background of the Invention

In relational databases, data is organized into tables. A collection of such tables in
a database is referred to as a table space. Database users sometimes inadvertently delete,
20 or drop, a table from a table space. Typically, a database management system (DBMS)
does not permit undelete of the drop action: once the table drop is committed, the table's
data is permanently deleted and cannot be brought back by way of such a command as an
undelete of the drop statement. Instead, the data must be restored from a backup and then
the data rolled forward by replaying stored transactions on the data, a potentially slow
25 process.

In certain DBMS environments, such as DB2**, recovering a dropped table is
made even more difficult by the fact that a table space restore followed by a roll forward
of the table space to a point in time prior to the drop cannot be done. This restriction
means that to roll forward to recover a dropped table, the entire database, and not only the

5 effected table space, must be rolled forward. This restriction on the roll forward of the table space is due to the minimum recovery time property. A table space must be rolled forward to at least the minimum recovery time so that it is synchronized with the information in the system catalog tables. The minimum recovery time is updated when data definition language (DDL) statements are executed against a table space, or against tables in a table space.

10 The minimum recovery time will be later than the time at which the table was dropped. Because of this fact, it is a requirement in typical DBMS environments that the user must recover the entire database. This means that the entire database becomes unavailable to other users while the recovery and rollforward of the database is being carried out. It is typically slower to perform the recovery and restore on the entire database than a recovery and rollforward on the effected table space, only.

15 The only means in which a dropped table can be currently recovered in many relational databases, such as DB2, is through a database restore followed by a database roll forward to a point in time just prior to the table drop. As indicated above, such an approach will make the database as a whole unavailable to users. It may also be difficult to pinpoint when a table was dropped and therefore data will often be inaccurately retrieved due to uncertainties about when the table in question was dropped. In addition, the structure of the table may no longer be accurately known.

20 It is therefore desirable to have a DBMS in which tables that have been dropped may be recovered without the need to recover and roll forward the entire database and which permits the point at which the table was dropped to be determined with some accuracy, as well as to determine the structure of the table at the time of the drop.

25 SUMMARY OF THE INVENTION

A method, system, and program for recovering a dropped table are provided. One or more table spaces are specified prior to one or more tables being dropped from the specified one or more table spaces. When a table is to be dropped from the one or more specified table spaces, a table identifier, a time stamp, and table definition attributes are

stored for the dropped table in a data structure and the table is dropped. Upon receiving a request to restore a table space from the one or more specified table spaces, each of the one or more tables in the table space being restored is recovered using the data structure.

Advantages of the preferred embodiments include the ability to restore and
5 rollforward the table space of the dropped table without having to restore the entire database. In addition, the dropped table history records a timestamp for the drop of the dropped table to permit the rollforward to the drop of the table to be carried out with some accuracy. The structure of the table which has been dropped is also available for use in the recovery of the table.

10

BRIEF DESCRIPTION OF THE DRAWINGS

Referring now to the drawings in which like reference numbers represents corresponding parts throughout:

Figure 1 is a block diagram showing the processing of the table space to recover a
15 dropped table, in accordance with the system of the preferred embodiments.

In the drawing, the preferred embodiments of the invention are illustrated by way of example. It is to be expressly understood that the description and drawing are only for the purpose of illustration and as an aid to understanding, and are not intended as a definition of the limits of the invention.

20

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

Referring to Figure 1, there is a block diagram showing the processing of a table space to recover a dropped table. Pre-drop table space 10 represents a table space in a database which pre-drop table space 10 is in a state immediately preceding the point of
25 commitment for drop table command 12. In the preferred embodiments, the database is a relational database which implements SQL, such as the DB2 DBMS. The system of the preferred embodiments permits a user to turn "on" a dropped table recovery flag. A sample SQL command to turn "on" the dropped table recovery flag for an example table space ts1 is:

alter table space ts1 using dropped table recovery on

If this flag is turned on for a table space, the system of the preferred embodiments will create an entry in a dropped table history 14 data structure, when a table (ts1, in the above example) is dropped from pre-drop table space 10. Dropped table history 14 records a dropped table ID, which is a unique identifier for the dropped table. Dropped table history 14 also provides a timestamp which reflects the time at which the dropped table was dropped from the table space, as well as information about the structure of the table (table definition attributes). The option to record an entry in the dropped table history 14 is table space specific. In the preferred embodiments, the flag may be queried in the "syscat.tablespace" catalog table. There is a drop_recovery column which may be queried. The flag may be turned "off" and "on" by the user as required. When a table is dropped in the system of the preferred embodiments, the status of the flag is determined by the system and if the flag is "on", then information about the dropped table is stored in dropped table history 14.

In the system of the preferred embodiments, the dropped table history 14 may be accessed by the LIST HISTORY DROPPED TABLE command. This command returns the dropped table ID, timestamp of the drop, and information about the structure of the table. An example of the command for the database test_db is as follows:

list history dropped table all for test-db

After accessing the dropped table history 14 in this way, the user of the system of the preferred embodiments may then restore the table space. An example of such a command for table space ts1 in database test_db is:

restore db test_db tablespace (ts1)

In Figure 1, a post-drop table space 16 represents the state of the table space after the drop table command 12 has been committed. It is expected that a number of

transactions (not shown) will have been processed and that the post-drop table space 16 may differ significantly from the state of the pre-drop table space 10.

Figure 1 shows the execution of restore command 18. This results in a restored table space 20. The restored table space 20 must pre-date the pre-drop table space 10.

5 The user then uses the rollforward command in the system of the preferred embodiments to replay stored transactions 22 against restored table space 20. The system of the preferred embodiment permits the user to stop the rollforward of restored table space 20 at the timestamp found in the dropped table history 14 for the table which was inadvertently dropped. The rollforward command of the preferred embodiment permits
10 the user to specify that the rollforward is to stop at the drop point for a table with a given dropped table ID. In Figure 1, transactions 22 are shown as being played against the restored table space 20, up to the point indicated by dropped table history 14. The result is the pre-drop table space 24. The state of the pre-drop table space 24 is intended to be materially equivalent to the state of pre-drop table space 10. The system has effectively
15 permitted the recreation of the table space to the state of pre-drop table space 10, to permit the data from the dropped table to be extracted from the table space as it existed immediately prior to the drop table command 12 being committed in the system. Once the appropriate transactions have been replayed against restored table space 20 to create pre-drop table space 24, a copy of the dropped table is made in flat file 26.

20 To accomplish this the user specifies a file location and the table data as it existed prior to drop will be written to the file as a flat file with ASCII delimiters. An example rollforward command for the database test_db, table space tsl, dropped table ID "00000000000000b60000" is shown where the flat file representation of the table is to be stored in file location /temp/ffile.

25 rollforward db test_db to end of logs and stop
tablespace (tsl) recover dropped table
00000000000000b60000 to temp/ffile

With reference to Figure 1, once the data from the dropped table in pre-drop table space 24 has been saved to flat file 26, further transactions 28 may be replayed on pre-drop table space 24 to result in recovered table space 30. The dropped table may be redefined in recovered table space 30 by create table command 32. The result is
5 intermediate recovered table space 34. Load data command 36 repopulates the table created by create table command 32 from flat file 26 to create recovered table in current table space 38. In this way, the dropped table is recovered in the table space without the need to restore the entire database. It will be appreciated by those skilled in the art that although in Figure 1 the various table spaces 10, 16, 20, 24, 30, 34 and 38 are shown as
10 separate items, a typical implementation of the preferred embodiment uses the same table space for all items, although the state of the table space varies over time, as described above.

By replaying transactions on the table space only, rather than on the database as a whole, the database remains available for use by other users and the time to recover the
15 dropped table is potentially shorter. The data from the dropped table is stored in flat file 26 to permit the data to be recovered into a table space which has different attributes than the table space from which the table was dropped. Dropped table history 14 includes dropped table characteristics which permit the table to be redefined in intermediate recovered table space 34 by create table command 32.

20 In SQL commands, the above procedure may be carried out using the following steps. An example of how to define the table into which the dropped table data will be copied is the following SQL command:

```
create table "test ." table1 " ( "col1" integer ) in "tsl"
```

The table is then repopulated from flat file 25 using a command:

25 load from data of del insert into table1

As will be apparent to one skilled in the art, there are other variations on the above approach to recovering the dropped table, given the system of the preferred embodiment. In particular, the recovery of the dropped table may be accomplished by restoring the database as a whole into a different system than the first database. The table

space can then be restored in the second system and the rollforward and recovery of the dropped table carried out "offline". This will permit the existing database to be used without making even the table space of the dropped table unavailable to users of the system. Another alternative to the steps described above is to make a copy of the current table space before carrying out the restore of the table space. Once flat file 26 is created, transactions 28 may be ignored and recovered table space 30 may be created by restoring the copy of the current table space, rather than by replaying transactions 28 on pre-drop table space 24. This alternative approach is advantageous where there are significant numbers of transactions in transactions 28.

As will also be apparent to those skilled in the art, flat file 26 may have other uses, and the data in flat file 26 may be exported to other applications or used in the database in other contexts than the repopulation of the dropped table.

Although preferred embodiments have been described here in detail, it will be appreciated by those skilled in the art, that variations may be made thereto, without departing from the spirit of the invention or the scope of the appended claims.

****DB2 is a registered trademark of International Business Machines Corp.**